

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization International Bureau



(43) International Publication Date
11 November 2004 (11.11.2004)

PCT

(10) International Publication Number
WO 2004/098109 A2

(51) International Patent Classification⁷:

H04L

(74) Agents: BRANCH, John, W. et al.; Darby & Darby P.C.,
P.O. Box 5257, New York, NY 10150-5257 (US).

(21) International Application Number:

PCT/IB2004/000765

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CI, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PI, PT, RO, RU, SC, SD, SE, SG, SK, SI, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(22) International Filing Date: 16 March 2004 (16.03.2004)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

60/467,066 30 April 2003 (30.04.2003) US
10/635,205 5 August 2003 (05.08.2003) US

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(71) Applicant (for all designated States except US): NOKIA CORPORATION [FI/IT]; Keilalahdentie 4, ITN-02150 Espoo (FI).

(71) Applicant (for LC only): NOKIA, INC. [US/US]; 6000 Connection Drive, Irving, TX 75039 (US).

(72) Inventors; and

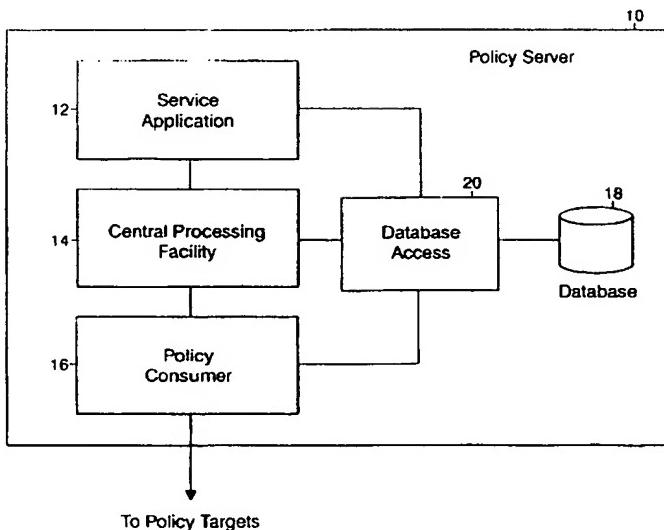
(75) Inventors/Applicants (for US only): LIONG, Yin, Ling [-/US]; 64 Mill St. #10, Woburn, MA 01801 (US). BARNES, Roberto [PA/FI]; Sateenkaari 3 D68, FIN-02100 Espoo (FI). LI, Man [CA/US]; 3 Overlook Drive, Bedford, MA 01730 (US).

Published:

— without international search report and to be republished upon receipt of that report

{Continued on next page}

(54) Title: SYSTEM, APPARATUS AND METHOD FOR SUPPORTING CONSTRAINT BASED ROUTING FOR MULTI-PROTOCOL LABEL SWITCHING TRAFFIC ENGINEERING IN POLICY-BASED MANAGEMENT



(57) Abstract: A policy server is arranged to construct device-neutral policies for configuring CBR for MPLS traffic engineering across a network. The policy server translates the device-neutral policies into device-specific commands (e.g. link attribute definitions and affinity profiles). The policy server defines the link attributes, assigns the link attributes to network interfaces, establishes affinity profiles and attaches the affinity profiles to MPLS tunnels. The link attribute definitions and affinity profiles are shared across the network to construct the policies such that IP operators can configure CBR easily across a network.

WO 2004/098109 A2

WO 2004/098109 A2



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

**SYSTEM, APPARATUS AND METHOD FOR SUPPORTING CONSTRAINT
BASED ROUTING FOR MULTI-PROTOCOL LABEL SWITCHING TRAFFIC
ENGINEERING IN POLICY-BASED MANAGEMENT**

5

Related Application

This utility application claims benefit under 35 United States Code § 119(e) of United States Provisional Application No. 60/467,066 filed on April 30, 2003.

Field of the Invention

10 This invention relates to data packet routing, and in particular, to policy management for configuring constraint based routing in multi-protocol label switching traffic engineering.

Background of the Invention

15 Packet-forwarding policies are in place to administer, manage and control access to network resources. Policy-based management employs a policy server to manage the network as a whole. The policy server translates business goals or policies into configurations of network resources and automates the configurations across multiple different network elements and different technologies (e.g. MPLS and DiffServ). The centralized approach ensures policy consistency across multiple network elements.

20 Multi-protocol label switching (MPLS) is a packet-forwarding technology that gives internet protocol (IP) operators a high degree of control over the paths taken by packets on their networks. MPLS may be used for traffic engineering purposes. Interior gateway protocols (IGP), such as open shortest path first (OSPF) and intra-domain intermediate system to intermediate system routing protocol (IS-IS), route 25 IP packets based only on the destination address and the shortest path to reach the destination. In contrast, MPLS traffic engineering allows administrators to establish routes for certain customers based on information other than the shortest path, such as delay and bandwidth available along the path. Therefore, MPLS can relieve congestion

and maximize bandwidth utilization by allowing multiple paths between source and destination.

- Constraint based routing (CBR) is an example of MPLS traffic engineering. The operator does not specify the path explicitly but depends on a CBR mechanism that has been implemented in the network to determine the path. With CBR, every router advertises traffic-engineering attributes (e.g., maximum bandwidth, unreserved bandwidth, etc.) for its interfaces to all other routers. OSPF and IS-IS protocols have been extended for that purpose. As a result of the advertisement, every router obtains a traffic-engineering database in addition to the regular routing database.
- 5 When a label switch path (LSP) with QoS requirements needs to be established, the ingress router computes the optimal path based on the traffic-engineering database and a path selection algorithm. The router then signals the path establishment with a label distribution protocol. The explicit path selected is conveyed through the signaling protocol.
- 10

15

Summary of the Invention

This summary of the invention section is intended to introduce the reader to aspects of the invention. Particular aspects of the invention are pointed out in other sections herein below, and the invention is set forth in the appended claims, which alone demarcate its scope.

20

The present invention is directed to a system for configuring constraint based routing (CBR) in multi-protocol label switching (MPLS) traffic engineering with a policy-based management approach across a network. The system includes network interfaces where link attribute definitions are set, affinity profiles that relate to preferred link attributes, and a policy server. The policy server attaches the affinity profiles to a

25 MPLS tunnel such that the affinity profiles are shared across the network to construct service and network policies.

25

Another aspect of the present invention is directed to an apparatus for configuring CBR in MPLS traffic engineering with a policy-based management approach across a network. The apparatus includes a service application, a central processing facility and a policy consumer. The service application configures policies.

30

The central processing facility translates the policies into device-neutral policy parameters. The policy consumer translates the device-neutral policies into device-specific commands and deploys the device-specific commands to policy targets, such that the policies are constructed across the network.

5 Another aspect of the present invention is directed to a method for supporting CBR for MPLS traffic engineering across a network. The method includes: defining link attributes; assigning the link attributes to network interfaces; establishing affinity profiles that specify preferred traffic engineering attributes; and attaching the affinity profiles to MPLS tunnels such that service and network policies are constructed
10 across the network.

Another aspect of the present invention is directed to an apparatus for configuring CBR in MPLS traffic engineering with a policy-based management approach across a network. The apparatus comprises: a means for defining link attributes; a means for assigning the link attributes to network interfaces; a means for 15 establishing affinity profiles that specify preferred traffic engineering attributes; and a means for attaching the affinity profiles to MPLS tunnels such that service and network policies are constructed across the network.

Brief Description of the Drawings

20 FIGURE 1 is a block diagram illustrating the architecture of a policy server;

FIGURE 2 is a block diagram illustrating the architecture of a service application;

25 FIGURE 3 is a block diagram illustrating the interaction of a service policies object with the service application;

FIGURE 4 is a block diagram illustrating the structure of a network policies object;

FIGURE 5 is a block diagram illustrating the policy elements of a device object;

FIGURE 6 is a block diagram illustrating the policy elements of a MPLS tunnels object;

FIGURE 7 is a diagram illustrating a network implementing the policy server of the present invention;

5 FIGURE 8 is a flow diagram illustrating a method for configuring link attributes and affinity profiles to support CBR for MPLS traffic engineering;

FIGURE 9 is a flow diagram illustrating a method for deploying a policy to policy targets; and

FIGURE 10 is a block diagram illustrating the method implemented in a
10 3G network, in accordance with aspects of the invention.

Detailed Description of the Preferred Embodiment

In the following detailed description of exemplary embodiments of the invention, reference is made to the accompanied drawings, which form a part hereof,
15 and which is shown by way of illustration, specific exemplary embodiments of which the invention may be practiced. These embodiments are described in sufficient detail to enable those skilled in the art to practice the invention, and it is to be understood that other embodiments may be utilized, and other changes may be made, without departing from the spirit or scope of the present invention. The following detailed description is,
20 therefore, not to be taken in a limiting sense, and the scope of the present invention is defined only by the appended claims.

Briefly stated, a policy server is arranged to construct device-neutral policies for configuring constraint based routing (CBR) for multi-protocol label switching (MPLS) traffic engineering across a network. The policy server translates the
25 device-neutral policies into device-specific commands. The policy server defines link attributes, assigns the link attributes to network interfaces, establishes affinity profiles and attaches the affinity profiles to MPLS tunnels. The link attribute definitions and the affinity profiles are shared across the network to construct the policies such that IP operators can configure CBR easily across a network.

FIGURE 1 illustrates a block diagram of the architecture of a policy server that supports differentiated services (Diffserv) over MPLS traffic engineering. Policy server 10 includes service application 12, central processing facility 14, and policy consumer 16 all coupled to database 18 through database access 20 which 5 provides interfaces for database 18 read and write operations.

Service application 12 is a graphical user interface (GUI) that allows IP operators to configure policies (i.e., operators can add, delete or modify policies.) Central processing facility 14 translates the policies into device-neutral policy parameters and stores the policy parameters in database 18. Central processing facility 10 14 also conducts policy verification, conflict detection and resolution. Policy consumer 16 provides an interface for central processing facility 14 to communicate with policy targets. A policy target refers to any network node where the policy is enforced such as a router. Policy consumer 16 translates the device-neutral policies stored in database 18 into device-specific commands, and deploys the device-specific commands to the policy 15 targets.

According to one embodiment, service policies and network policies are used. Service policies refer to rules that govern the treatment of individual customer traffic, such as traffic classification, metering and marking. Network policies refer to rules that govern the treatment of aggregated traffic, such as queue and scheduler 20 configurations for behavior aggregates. For example, an IP operator may modify service policies when a new service level agreement (SLA) is created or when an existing SLA is modified. An IP operator may modify network policies when new facilities are added into a network or when traffic patterns are changed significantly. In a large network, the same set of policies can be applied to multiple routers or network 25 interfaces. To achieve better scalability, network interfaces are assigned role names and the policies are specified and associated with the roles name. Network interfaces identified by the same role names receive the same set of policies.

A common information model describes policies as relationships between network objects. The information model derives database schema for database 30 18 such that the policy stored in database 18 is device-neutral. In a multi-vendor

environment where routers from different vendors are configured through different protocols, policy server 10 translates the policy stored in database 18 into device-specific commands and delivers the resulting policy through the specific protocol. The information model enables a consistent provision of policies across multi-vendor
5 networks.

FIGURE 2 illustrates a block diagram of the architecture of service application 12. Service application 12 includes services object 24, application object 26, customer object 28, devices object 30, MPLS tunnels object 32 and network policies object 34. Services object 24, applications object 26, and customer object 28 are related
10 to service policies. Devices object 30, MPLS tunnels object 32, and network policies object 34 are related to network policies.

Services object 24 allows IP operators to define the services to be provided to customers. According to one embodiment, each service is defined as a mapping between a service name and any of the fourteen DSCP numbers defined by the
15 Internet Engineering Task Force (IETF). Application object 26 allows IP operators to define applications for the classification of traffic flows such as protocol, source port and destination port ranges. Customer object 28 allows IP operators to define service policies for each customer. Customer object 28 may also contain host groups and metering profiles associated with each customer. Devices object 30 contains
20 information about network elements that receive the policies from policy server 10. IP operators use devices object 30 to assign role names to the network interfaces. Most other policy target information in devices object 30 is directly imported into database 18 from an element management database. MPLS tunnels object 32 allows IP operators to configure the MPLS network components such as tunnel group, explicit routes, CBR
25 and EXP-PHB mapping, i.e., MPLS tunnels object 32 specifies the MPLS and DiffServ/MPLS policies. Network policies object 34 allows IP operators to define the network policies that specify the queuing and scheduling treatment of a specific service and that are attached to specific network interfaces.

FIGURE 3 illustrates a block diagram of the interaction between a
30 service policies object and service application 12. Service policies object 36 includes

rules that govern the treatment of individual customer traffic such as classification and metering rules. The classification rules include an application name supplied by application object 26 and source/destination host groups supplied by customer object 28. The metering rules use an algorithm specified by traffic profiles which are supplied 5 to service policies object 36 by customer object 28. Devices object 30 provides role names to service policies object 36 such that service policies object 36 can be applied to the network interfaces. MPLS tunnels object 32 provides information about tunnel group and tunneling mode to service policies object 36. Service policies object 36 defines non-conformance action (e.g., marking or dropping) after metering, and maps 10 the traffic into the desired MPLS tunnels and tunneling mode.

FIGURE 4 illustrates a block diagram of the structure of the network policies object. Devices object 30 allows operators to assign role names to the network interfaces. Network policies object 34 specifies the Diffserv configuration and applies the configuration to the role names. Network policies object 34 then deploys the 15 network policies to the network interfaces through the associated role names.

Services object 24 provides service names to network policies object 34. The service name selects the Diffserv class that supports classified traffic. Network policies object 34 defines the PHB used to support a particular service or Diffserv class, such as queuing and scheduling. In each network interface, all LSPs share the same 20 queue if they have the same PSC, and no per-LSP queuing is employed.

FIGURE 5 illustrates a block diagram of the policy elements of the device object. The policy elements set link attributes 37 in each network interface 38. Policy server 10 defines the semantics of link attributes 37 and affinity profiles supported by link attribute definition object 40. Link attribute definition object 40 25 allows IP operators to define a high-level representation for each affinity bit, such as the specific data rate, link technology or virtual private network (VPN) membership. In one embodiment, each representation consists of a bit position, attribute names and a category. In another embodiment, each representation consists of a bit position and an attribute name. An example of an attribute name is "Ethernet" and an example of a 30 category is "Link Technology". Link attributes 37 are then stored in database 18 and

shared across different objects. For instance, link attributes 37 can be added to each network interface. Thus, IP operators can select the correct link attributes to describe the network interface.

- FIGURE 6 illustrates a block diagram of the policy elements of the
- 5 MPLS tunnels object. The policy elements include tunnel group 42, tunnel characteristics object 44, explicit route object 46, affinity profile object 48, and EXP-service map 50. Tunnel group 42 is a set of tunnels that share the same properties and form a certain topology, such as a mesh or a star. Tunnel group 42 configures (i.e., creates, maintains and deletes) MPLS tunnels such that IP operators do not need to
- 10 create tunnels individually. Instead, IP operators specify the end-point routers and the inter-connecting topology. The internal routes traversed by a tunnel are determined by the underlying routing protocols, such as OSPF or IS-IS. In the case of star topology, the hub and the spokes are specified explicitly.

Tunnel characteristics object 44 includes profiles that specify the

15 parameters for resource reservation and policing per MPLS tunnel. Explicit route object 46 allows IP operators to explicitly specify the optimal path for forwarding packets. Explicit route object 46 contains a list of interface links that an explicit route LSP traverses through. The list of IP addresses can be a partial or complete route. There can be more than one explicit route list and thus a pointer is necessary to link the

20 list to tunnel group 42. In addition, each explicit route LSP is unidirectional. Therefore, a bidirectional LSP requires configuration of two explicit route LSPs in opposite directions.

There are two steps in configuring tunnel group 42 using CBR: setting the link attributes and attaching affinity profile object 48 to tunnel group 42. Affinity profile object 48 defines affinity profiles that select the preferred link attributes in routing. An affinity profile indicates whether each attribute defined in link attribute definition object 40 is preferred, not preferred, or "don't care." An affinity profile is linked explicitly to a tunnel group. Central processing facility 14 then converts the affinity profile into device neutral information for storage in database 18, which is then

25 used with an extended IGP protocol to support CBR.

- The linking of tunnel group 42 in service policies object 36 identifies the LSPs that carry customer traffic. Tunneling mode decides which Diffserv code point (DSCP) is carried in the IP headers when a packet exits the MPLS network. There are two essential modes of tunneling: pipe mode and uniform mode. For pipe mode, the 5 egress router keeps the DSCP of the encapsulated IP header. For uniform mode, the egress router overwrites the original DSCP with the Diffserv information in the MPLS network. Tunnel group 42 and tunneling mode are used if operators use MPLS to carry customer traffic. If MPLS is not used, tunnel group 42 and tunneling mode can be left empty.
- 10 FIGURE 7 illustrates a diagram of a network implementing the policy server according to aspects of the invention. Network 54 includes two edge routers 56 and two core routers 58. Policy server 10 can construct multiple LSPs between edge routers 56 through explicit route or CBR, as indicated by arrows 60. Policy server 10 can also map customer X and customer Y into one LSP or split them into two LSPs by 15 service policies. In addition, policy server 10 can assign different customer applications into different Diffserv classes by service policies, and allocate the resources accordingly by network policies.

According to embodiments of the invention, two approaches for policy deployment onto routers 56, 58 can be implemented. The first approach is generating a 20 new configuration file to replace the file currently in use by the router. FIGURE 8 illustrates a flow diagram of the second approach. The current configuration of a router is determined at block 62. The configuration includes information about access lists, classification rules, policies, MPLS tunnels, and the like. The router is reconfigured at block 63. The router is reconfigured by deleting unwanted policies and replacing the 25 unwanted policies with new command line interface (CLI) commands based on the current set of policies to be deployed.

FIGURE 9 illustrates a flow diagram of a method implemented in the policy server for configuring link attributes and affinity profiles according to aspects of the invention. An attribute editor is added to the policy server at block 64. The 30 attribute editor allows operators to define and edit the meaning of attributes supported

across the network. Examples of the definition can be related to geography, link speed, VPN group membership, and the like.

The link attributes defined in the attribute editor are configured for each network interface at block 66. Any change in the attribute editor is automatically reflected in the network interface configuration. The GUI for network interfaces displays the attribute names defined in the attribute editor thus allowing IP operators to select the link attributes that describe the characteristics of any network interface. Each interface can have different link attribute settings. Thus, the attribute editor hides the device-specific details for supporting link attributes and allows IP operators to 10 configure the link attributes easily.

Affinity profiles are added to the policy server at block 68. Each affinity profile specifies which link attributes are preferred, which attributes are not preferred and which attributes are "don't care" values. Any modification in the attribute editor is automatically reflected in the affinity profiles. IP operators can configure (i.e., create, 15 delete and edit) the affinity profiles. The policy server translates each affinity profile into device-specific commands at block 70. The policy server stores device-neutral information associated with each affinity profile in the policy server database at block 72.

An affinity profile is attached to an MPLS tunnel or tunnel group by the 20 policy server at block 74. A tunnel group refers to an identifier that represents a collection of MPLS tunnels that share the same properties and are connected by a certain topology. A tunnel group provides a mechanism to attach an affinity profile to multiple tunnels. An extended interior gateway protocol (IGP) supports MPLS traffic engineering and selects the optimal path based on affinity profiles and advertised link 25 attributes. The affinity profiles are shared across the network. Thus, two different tunnel groups or MPLS tunnels can reuse the same affinity profile. The policy server uses the affinity profile in constructing the device-specific commands to establish the MPLS tunnels.

The above-described methods can be implemented in an existing MPLS 30 policy server to support CBR for MPLS traffic engineering. FIGURE 10 illustrates an

example of the methods implemented in third generation (3G) networks. MPLS tunnels 82 are established across an IP transport network 84, which include public land mobile networks (PLMN), to provide VPN support. MPLS tunnels 82 provide pathways between IP networks 86 and radio access networks 88. The method provides IP

5 operators with a unified management tool that can configure serving edge routers 90 to support CBR. According to one embodiment, edge routers 90 include general packet radio service (GPRS) support nodes (SGSN), gateway GPRS support nodes (GGSNs) and third party MPLS routers.

The above specification, examples, and data provide a complete

10 description of the manufacture and use of the composition of the invention. Since many embodiments of the invention can be made without departing from the spirit and scope of the invention, the invention resides in the claims hereinafter appended.

WE CLAIM:

1. A system for configuring constraint based routing (CBR) in multi-protocol label switching (MPLS) traffic engineering with a policy-based management approach across a network, comprising:
 - 5 network interfaces that are arranged to have link attribute definitions set; affinity profiles that relate to preferred link attributes; and a policy server that is arranged to attach the affinity profiles to a MPLS tunnel such that the affinity profiles and the link attribute definitions are shared across the network to construct service and network policies.
- 10 2. The system of claim 1, wherein the affinity profiles are device-neutral and wherein the policy server is further arranged to translate the affinity profiles into device-specific commands.
- 15 3. The system of claim 1, wherein the policy server is further arranged to assign the link attributes to the network interfaces by determining the current configuration of a router, configuring the router, and issuing new commands based on the set of link attributes to be assigned.
- 20 4. The system of claim 1, wherein the policy server comprises a user interface that allows an operator to configure the service and network policies.
- 25 5. The system of claim 1, wherein the policy server is further arranged to configure the MPLS tunnels such that an operator can create an MPLS tunnel by specifying end-point routers and inter-connecting topology.
6. The system of claim 1, further comprising a database for storing the affinity profiles, the link attribute definitions and the link attributes of each network interface.

7. The system of claim 1, further comprising a link attribute definition object that is arranged to define the semantics of the link attributes and the affinity profiles.

5 8. The system of claim 7, wherein the link attribute definition object is further arranged to describe a policy target by defining a representation for each attribute bit and for each affinity bit in the affinity profile.

9. The system of claim 8, wherein the representation comprises at
10 least one of a bit position, an attribute name and a category.

10. An apparatus for configuring CBR in MPLS traffic engineering with a policy-based management approach across a network, comprising:
a service application that is arranged to have policies configured;
a central processing facility that is arranged to translate the policies into
15 device-neutral policy parameters; and
a policy consumer that is arranged to translate the device-neutral policies into device-specific commands, and that is further arranged to deploy the device-specific commands to policy targets, such that the policies are constructed across the network.

20

11. The apparatus of claim 10, further comprising a database for storing the policy parameters.

12. The apparatus of claim 10, wherein the service application
25 comprises a MPLS object that is arranged to configure MPLS tunnels such that an MPLS tunnel can be created by specifying end-point routers and inter-connecting topology.

13. The apparatus of claim 10, wherein the service application comprises a devices object and a network policies object, the devices object assigning role names to network interfaces, the network policies object deploying the policies to the policy targets based on the assigned role names such that the network interfaces 5 having the same role names receive the same policy.

14. The system of claim 10, wherein the service application further comprises a link attribute definition object that is arranged to define the semantics of link attributes, affinity profiles, and attribute bits such that a representation is defined 10 for each affinity bit in an affinity profile to choose a network interface.

15. The system of claim 14, wherein the representation comprises at least one of a bit position, an attribute name and a category.

16. A method for supporting CBR for MPLS traffic engineering 15 across a network, comprising:

defining link attributes;
assigning the link attributes to network interfaces;
establishing affinity profiles that specify preferred traffic engineering attributes; and
20 attaching the affinity profiles to MPLS tunnels such that service and network policies are constructed across the network.

17. The method of claim 16, further comprising storing the affinity profiles, the link attribute definitions and the link attributes of each network interface in 25 a database.

18. The method of claim 16, further comprising translating the affinity profiles into device-specific commands

19. The method of claim 16, further comprising storing device-neutral information associated with the affinity profiles in a database.

20. The method of claim 16, wherein assigning the link attributes to
5 network interfaces further comprises:
determining the command configuration of a router;
configuring the router; and
issuing new commands to the router based on the link attributes to be
assigned.

10 21. An apparatus for configuring CBR in MPLS traffic engineering
with a policy-based management approach across a network, comprising:
a means for defining link attributes;
a means for assigning the link attributes to network interfaces;
a means for establishing affinity profiles that specify preferred traffic
15 engineering attributes; and
a means for attaching the affinity profiles to MPLS tunnels such that
service and network policies are constructed across the network.

22. The method of claim 21, further comprising a means for
20 translating the affinity profiles into device-specific commands.

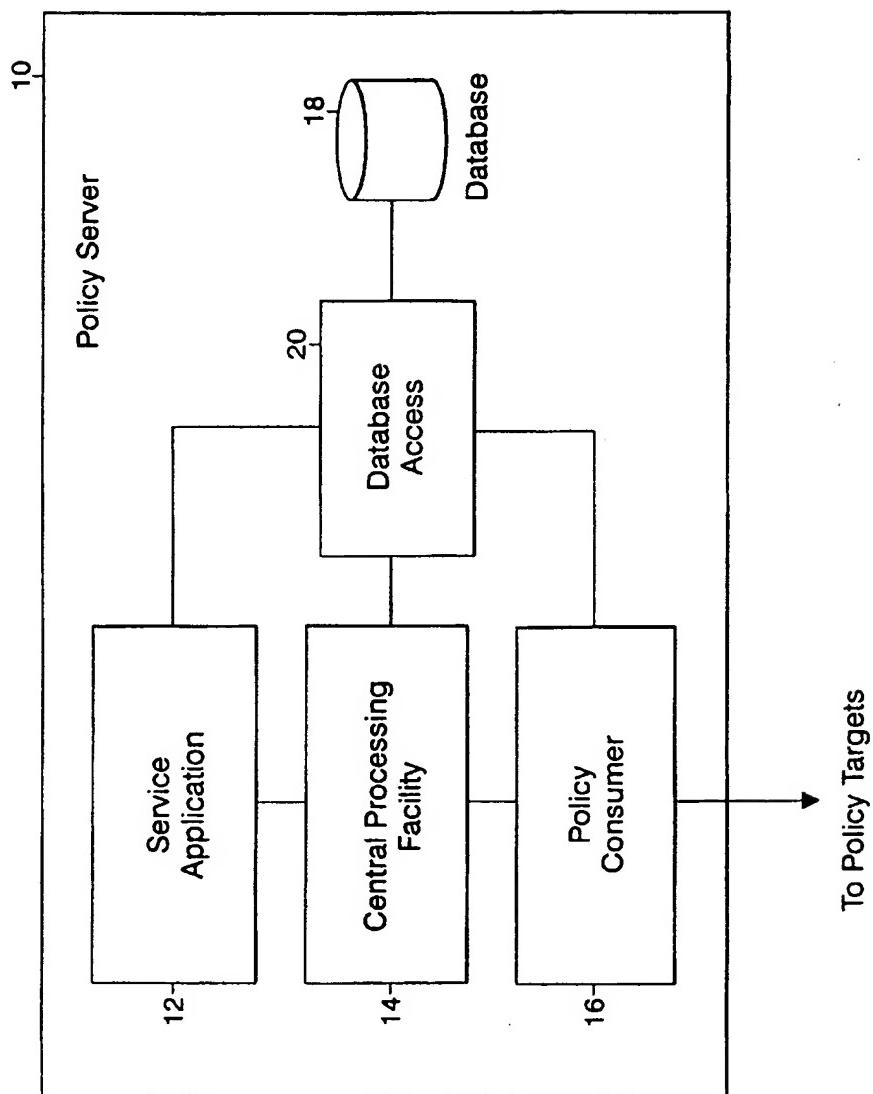
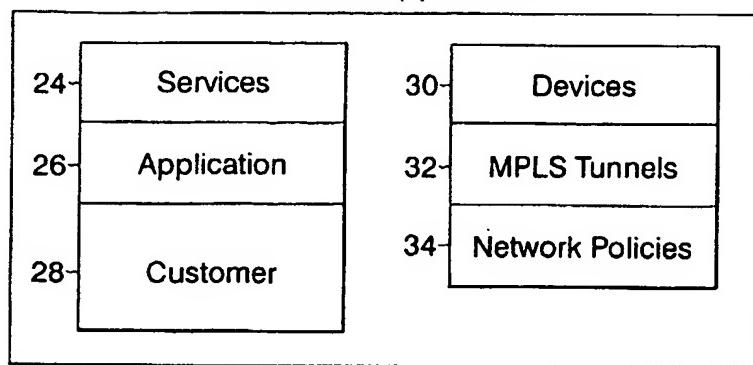
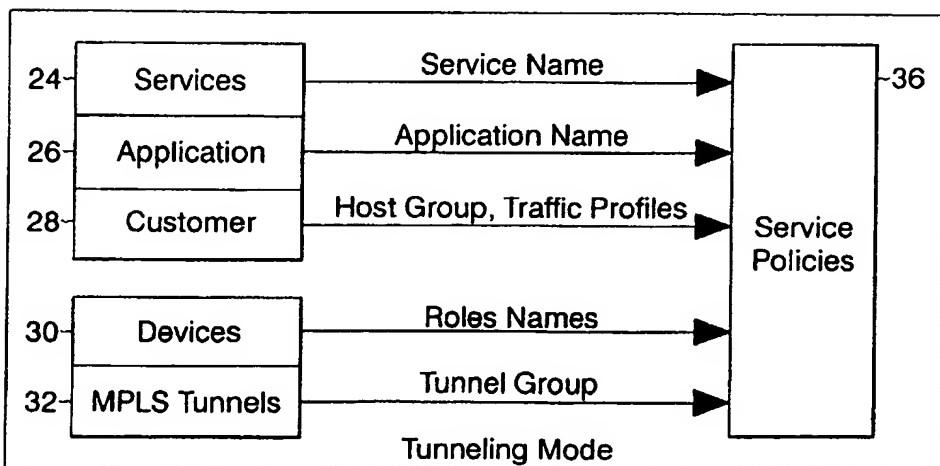
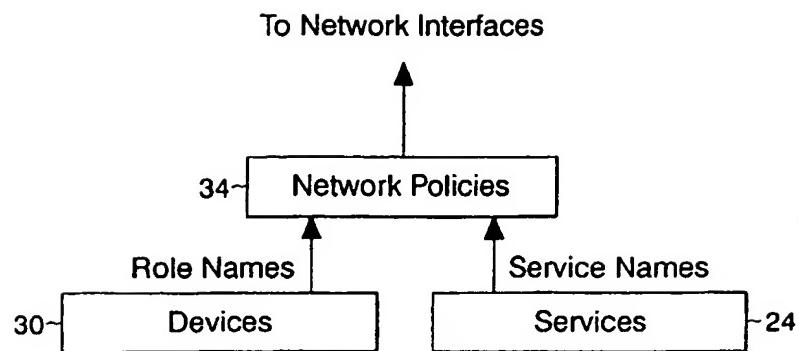
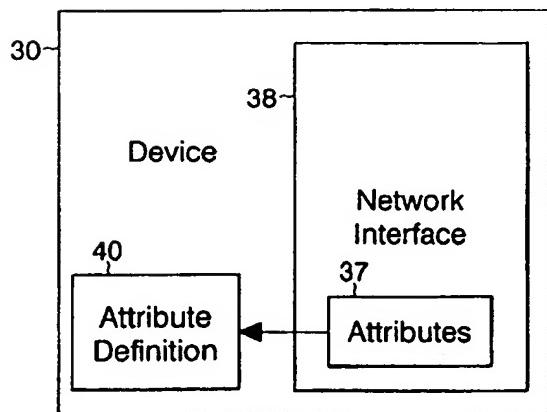


Figure 1

2/7

12**Service Application****Figure 2****Figure 3**

3/7

**Figure 4****Figure 5****SUBSTITUTE SHEET (RULE 26)**

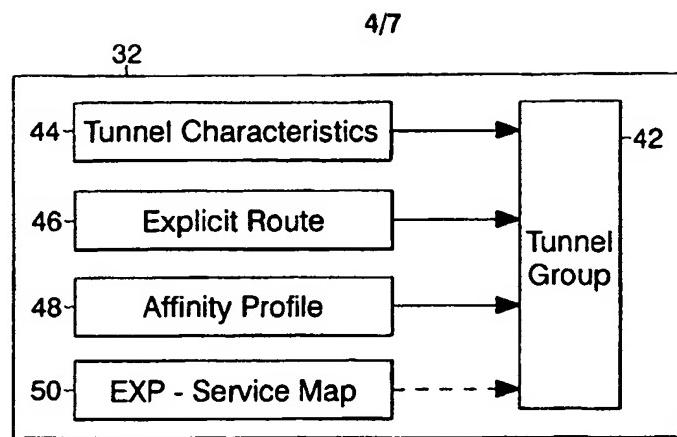


Figure 6

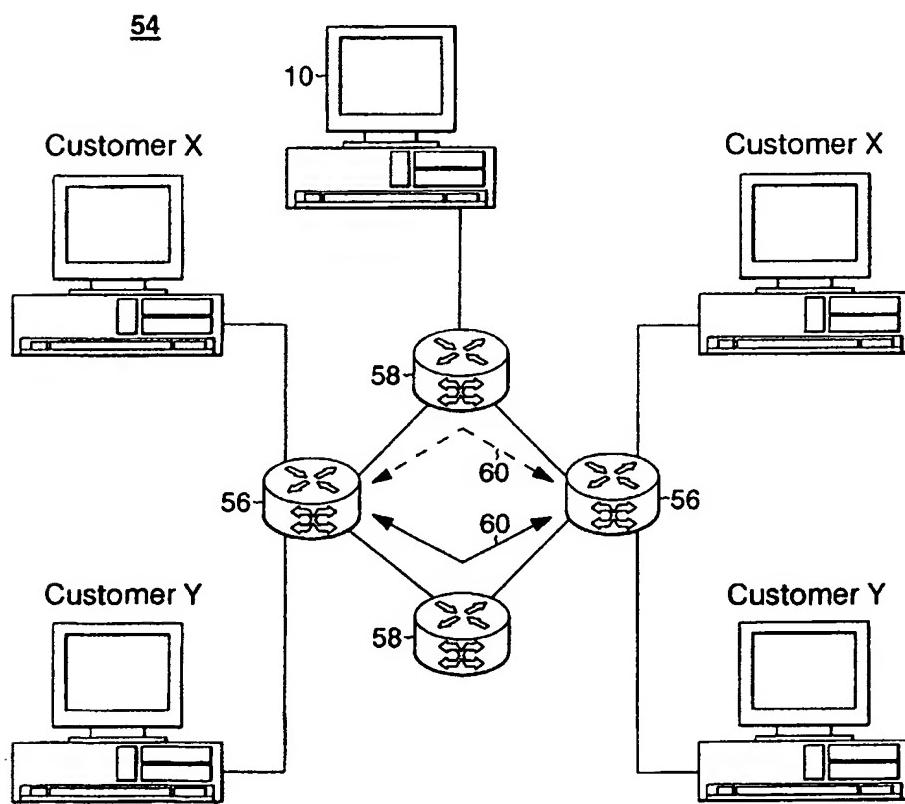


Figure 7
SUBSTITUTE SHEET (RULE 26)

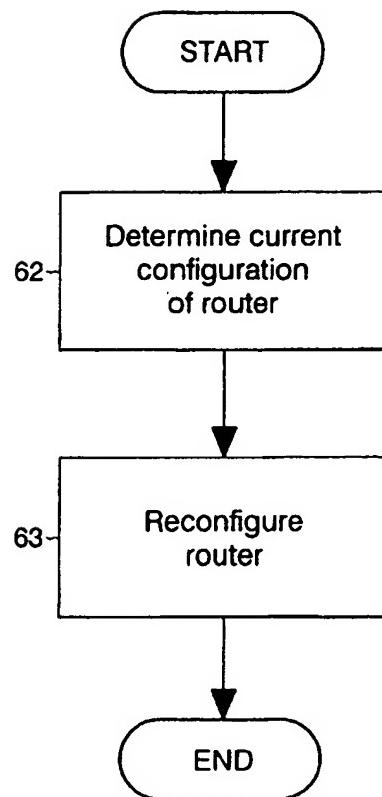
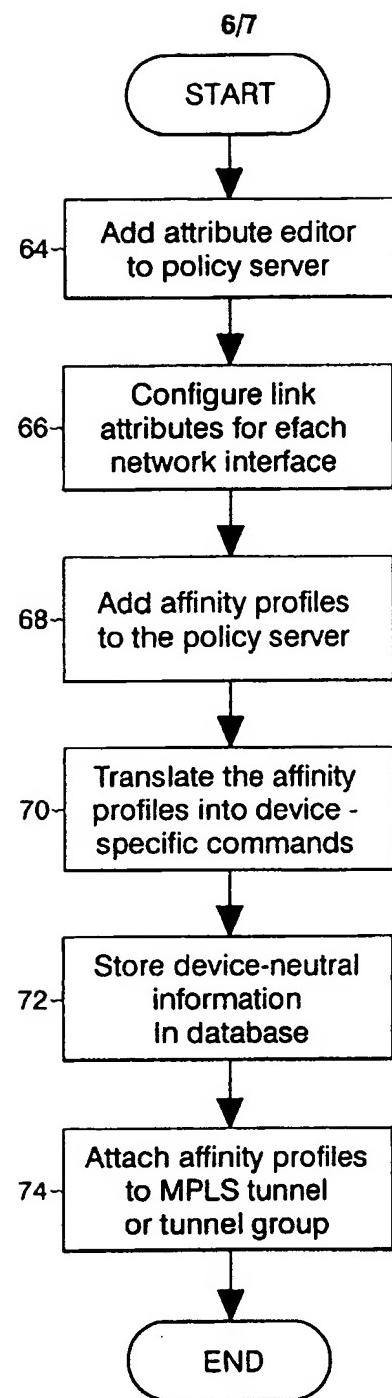


Figure 8

**Figure 9**

SUBSTITUTE SHEET (RULE 26)

7/7

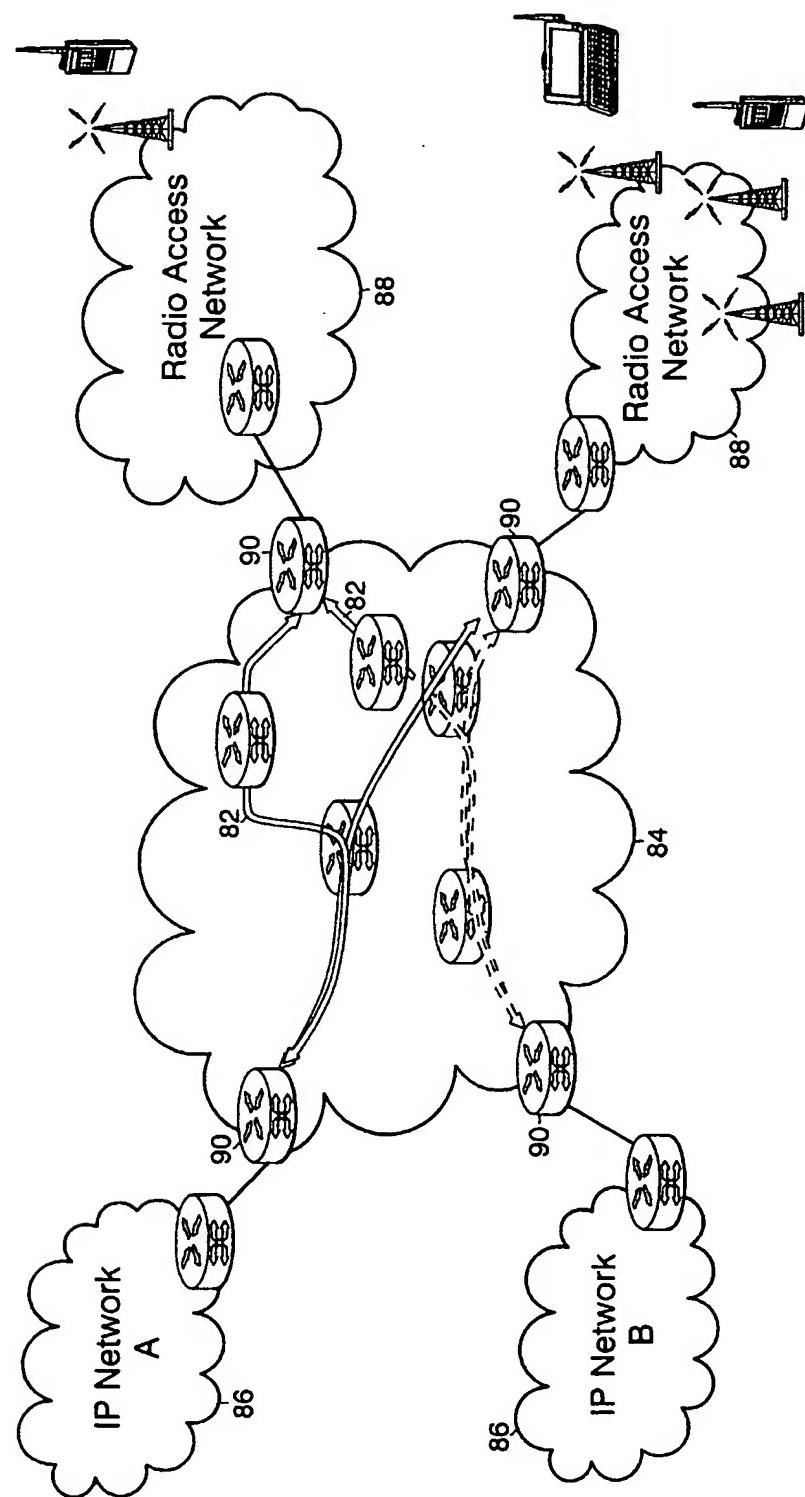


Figure 10

SUBSTITUTE SHEET (RULE 26)